4. CHARACTERISING SURFACES

§4.1. Holes

The Fundamental Theorem of Surfaces shows that every surface has the form $m\mathbf{D} + n\mathbf{P}$ or $m\mathbf{D} + n\mathbf{T}$ for some $m, n \geq 0$. Moreover the proof is constructive. If we're given a PIE we can follow the steps of the proof to

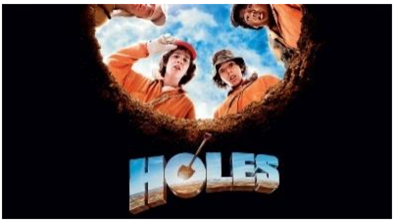

determine its structure in terms of disks, projective planes and toruses. However the surgery may be a little tedious. There's a better way! Instead we determine some invariants and use those to determine whether the surface is $m\mathbf{D} + n\mathbf{P}$ or $m\mathbf{D} + n\mathbf{T}$ as well as to determine the values of m and n.

We'll need 3 invariants altogether. The most obvious one is the number of holes in the surface. This will give us the number of disks in the decomposition.

A **hole** in a surface is a closed path in which every point is a boundary point (that is, has no neighbourhood homeomorphic to an open disk).

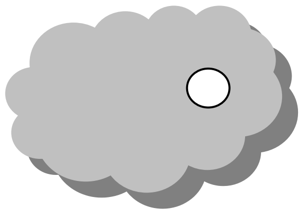

The number of holes is the number of disks, *m*, in the *m***D** + *n***P** or *m***D** + *n***T** decomposition. But the number of holes is not the same as the number of unidentified edges since it often happens that the boundary of a hole is made up of more than one such unidentified edge.

To find the number of holes (disks) in a surface:
(1) Use the identification of the edges to determine which vertices are identified.
(2) Then start at any vertex at the end of an unidentified edge and move to the vertex at the other end.
(3) Now find that vertex as an endpoint of another unidentified edge.
(4) Continue in this way until you have assembled sets of unidentified edges that collectively make up one closed curve (boundary of one hole).
(5) Continue in this way until all the unidentified edges have been accounted for.

Example 1: Find the number of holes in the following PIE.

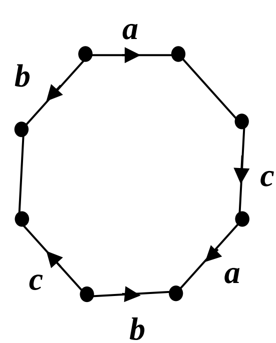

We begin by labelling a vertex adjacent to an unidentified edge as P.

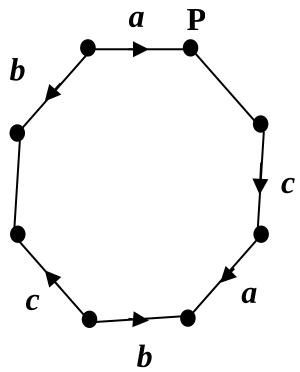

Since it's at the tip of an arrow marked ***a*** we must identify it with the tip of the other arrow marked as ***a***, and so we must label this also as P.

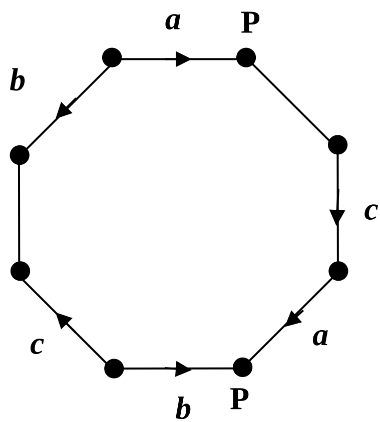

But this is also the tip of the arrow marked ***b*** and so the tip of the other arrow that's marked ***b*** must also be labelled as P.

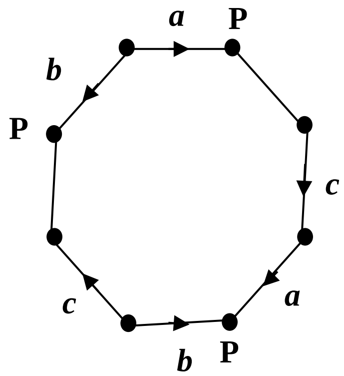

We can’t continue with P so we look for another vertex at the end of an unidentified edge and label it Q. We then continue the identification process and, in this case, we find that all vertices are labelled P or Q.

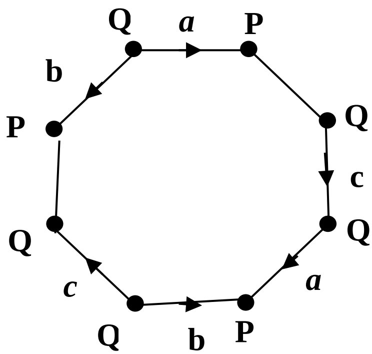

Taking into account the identified edges we find that we have 2 distinct vertices for this PIE. Now we have two unidentified edges but this doesn’t mean we have two holes. Notice that in this case we have an unidentified edge from P to Q. This isn’t a complete boundary. The other unidentified edge from P to Q (or from Q to P) makes up the other side of the boundary surrounding just one hole. This PIE has only one hole. Hence it is **D** + n**P** or **D** + n**T** for some $n \geq 0$. As to which of the two it is, and what is the value of n, we'll need to wait till a later section.

But suppose one of the two unidentified edges had endpoints P and P and the other had endpoints Q and Q. In that case these edges would have constituted two separate boundaries, surrounding two holes.

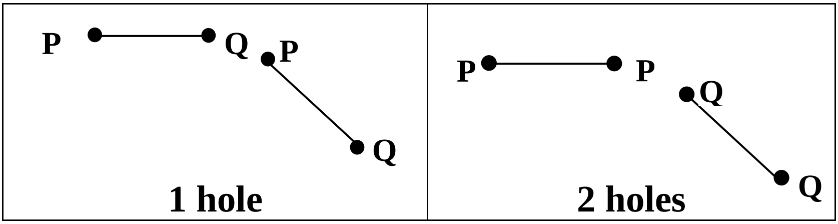

Had there been unlabelled vertices left we'd have continued with R etc. As long as we have a yet-to-be-considered unidentified edge we should proceed with a vertex at the end. But if no unidentified edges remain it doesn't matter where we go next. The reason for this is that strictly speaking we should go backwards as well as forwards from the vertex we start with each time. But if the vertex is at the end of an unidentified edge we can only proceed in one direction. And if there are no remaining unidentified edges the process of going in one direction will take us back to where we started.

Example 2: Consider the following PIE.

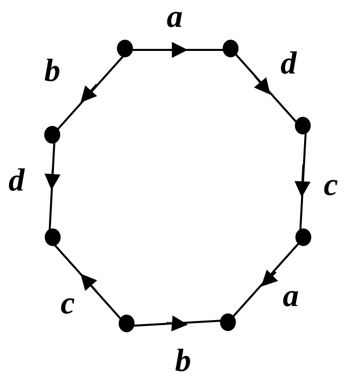

This one has no unidentified edges and so no holes. But for the next stage in the identification we still need to carry out the identifications of vertices.

We start anywhere and label it P. Then, going in either direction, we carry out the above process, identifying vertices. Eventually we will get back to our original vertex. All the vertices we encounter will be identified as P. If any vertices remain we repeat this whole process, identifying vertices as Q. We continue until all the vertices have been labelled.

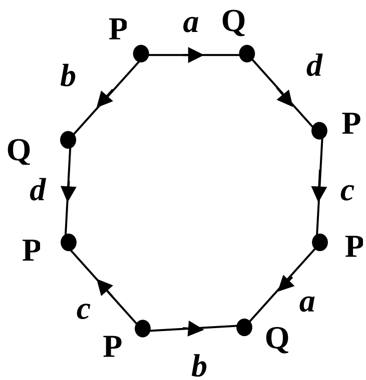

So this PIE has only 2 vertices and 4 edges.

§4.2. Orientability

We'll now answer the question of whether a given surface has the form $m\mathbf{D} + n\mathbf{P}$ or $m\mathbf{D} + n\mathbf{T}$. It's a question of 'orientability'.

On a surface, clockwise and anticlockwise orientations have no absolute meaning. To decide whether a rotation in a plane is clockwise or anti-clockwise we need to view it from outside the plane, in a third dimension. And even then it depends from which side of the surface we are viewing. Every clock runs anticlockwise if you can get inside it and view the hands from behind. But orientation does have a *relative* meaning on a surface.

Example 3: The arrows on this cylinder clearly have opposite orientations. But which is clockwise and which is anticlockwise depends on where we view it from.

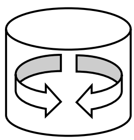

A **map** on a surface is a network of points (**vertices**) and lines (**edges**) joining them so that the surface is divided into regions (**faces**) such that the interior of each face is homeomorphic to an open disk.

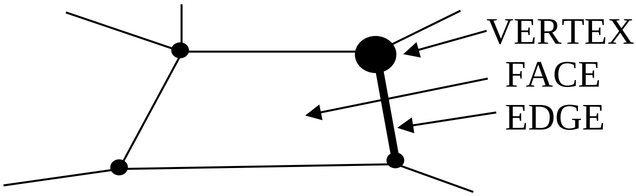

The reason for the terms ‘vertex’, ‘face’ and ‘edge’ comes from the fact that polyhedra (solid figures bounded by polygons) can be considered to be maps on a surface. The simplest polyhedra, the cube and the tetrahedron, are maps on a (topological) sphere.

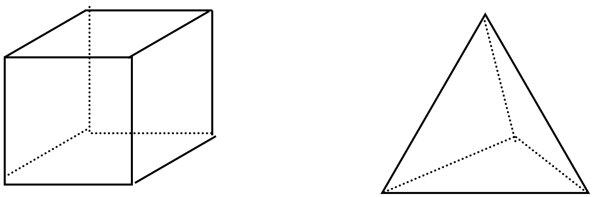

Cube **Tetrahedron**

An **oriented map** is a map in which each face is given an orientation in such a way that adjacent faces have the same orientation (that is, arrows on either side of the common border that follow the orientations are pointing in opposite directions).

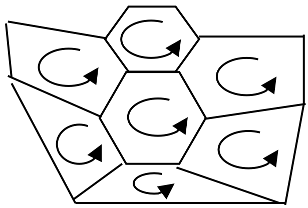

A surface is **orientable** if an oriented map can be drawn on it. Otherwise it is **non-orientable**.

Example 4: Spheres, disks, cylinders and toruses are orientable. Projective planes, Möbius Bands and Klein Bottles are non-orientable.

One way of thinking about orientability is to imagine a 2-dimensional white rabbit, like the one in *Alice in Wonderland*. He's carrying a big pocket-watch and runs about crying, "I'm late, I'm late, for a very important date". Suppose he has a twin rabbit. She also carries a pocket-watch, but she prefers staying in one place.

Is it possible that the first white rabbit could run about so far that when he returned to his sister he discovered that one of their watches was now running in the opposite direction to the other?

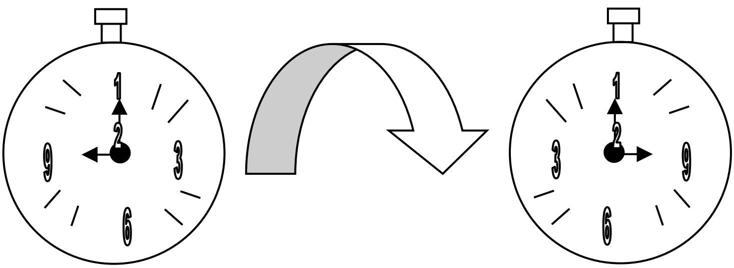

If they lived on a non-orientable surface this could happen. Suppose these rabbits lived on the surface of a

Möbius Band. Travelling right around the band the first rabbit's watch would have the opposite orientation to his sister's.

You may object that the first rabbit is now on the other side of the band. But you must remember that the Möbius Band is not really what you make out of a twisted strip of paper because paper only *appears* to be 2-dimensional. In fact it has thickness. The true Möbius Band is not the surface of the paper but rather the middle of the paper, half-way between the two sides. And the rabbits would be 2-dimensional rabbits, living within the Möbius Band rather than standing upon it.

Theorem 1: A surface that has a non-orientable subset is non-orientable.
Proof: If the surface is orientable a map can be drawn on it so that adjacent faces have the same orientation. Intersecting this with the subspace will give an orientable map on the subspace, which is a contradiction.

Theorem 2: Any PIE with a pair of like arrows is non-orientable.
Proof: Such a PIE contains a Möbius Band.

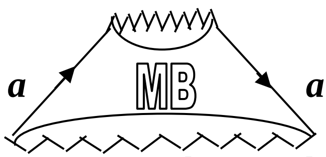

Example 5: A Projective Plane and a Klein Bottle are non-orientable.

Theorem 3: The sum of two orientable surfaces is orientable.

Proof: One simply draws an oriented map on each surface. Cut a hole inside one face in each surface and join the surfaces so that the boundaries of the holes have the same orientation. The resulting map is an orientable map on the combined surface.

§4.3. The Euler Characteristic

Every surface has the form $m\mathbf{D} + n\mathbf{P}$ or $m\mathbf{D} + n\mathbf{T}$ for some $m, n \geq 0$. It is orientability that decides whether it has the $m\mathbf{D} + n\mathbf{P}$ form or the $m\mathbf{D} + n\mathbf{T}$ form, and the number of holes gives us the value of m. It remains to determine n, the number of projective planes or toruses. We do this using an invariant called the **Euler characteristic**, which is ultimately a property of surfaces, but initially we define it for maps on surfaces.

The **Euler characteristic** ('Euler' is pronounced 'oiler') of a map with V vertices, E edges and F aces is:

$\boldsymbol{\chi = V + F - E}$.

Example 6: Maps on a disk:

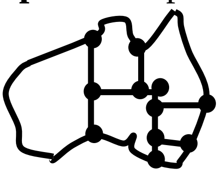

V = 11, F = 6, E = 16

$\therefore\ \chi = 11 + 6 - 16 = 1$.

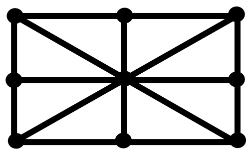

V = 9, F = 8, E = 16

$\therefore\ \chi = 9 + 8 - 16 = 1$.

Example 7: Maps on a sphere:

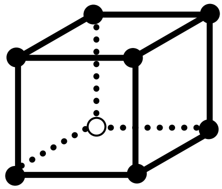

V = 8
F = 6
E = 12
$\therefore\ \chi = 8 + 6 - 12 = 2.$

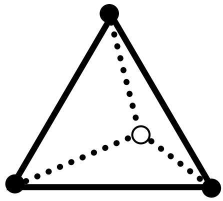

V = 4
F = 4
E = 6
$\therefore\ \chi = 5 + 5 - 8 = 2.$

Example 8: A map on a torus:

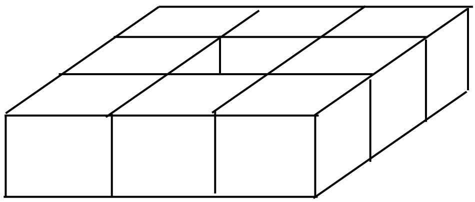

This consists of a 3 × 3 array of cubes with the middle one removed. Counting vertices, faces and edges is a little more difficult in this case because many of them are hidden. Adopt some sort of systematic approach and see if you get the same numbers as below. Don't forget to count the faces and edges in the middle.

V = 32
F = 32
E = 64
$\therefore\ \chi = 32 + 32 - 64 = 0.$

Notice that although the Euler characteristic of a map can vary, all maps on a given surface appear to have the same Euler characteristic. This observation is actually a fact, as we'll now see.

Theorem 4: Any two maps on the same surface have the same Euler characteristic.

Proof: The following operations leave the Euler characteristic of a map unchanged:

(1) Remove an edge separating two different faces and combine those two faces.

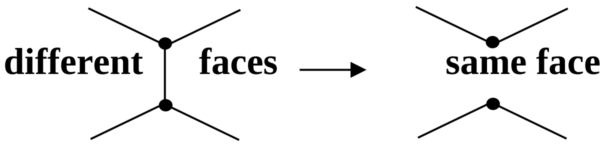

$$\left.\begin{aligned} V &\to V \\ F &\to F - 1 \\ E &\to E - 1 \end{aligned}\right\} \ \chi \to \chi.$$

(2) Remove a vertex of degree 1 (one that is an endpoint of only one edge) and remove the corresponding edge.

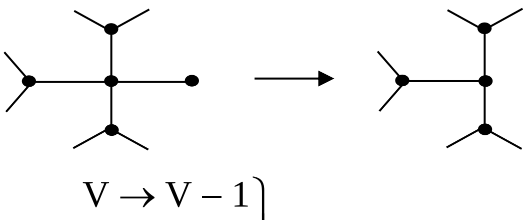

$$\left.\begin{aligned} V &\to V - 1 \\ F &\to F \\ E &\to E - 1 \end{aligned}\right\} \ \chi \to \chi.$$

(3) Remove a vertex of degree 2 (one that is an endpoint of exactly two edges) and combine those two edges into a single edge.

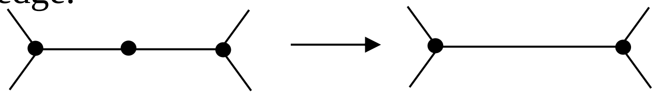

$$\left.\begin{array}{l} V \to V-1 \\ F \to F \\ E \to E-1 \end{array}\right\} \chi \to \chi$$

Suppose A and B are maps on a surface S. Superimpose the maps to form a composite map A + B. The map A + B can be changed back to either map A or map B by a sequence of operations (1), (2) and (3).

Hence $\chi(A) = \chi(A + B) = \chi(B)$.

Example 9:

A **B** **A + B**

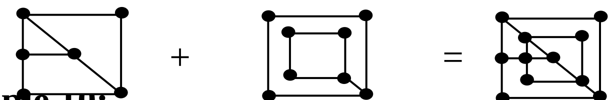

Example 10:

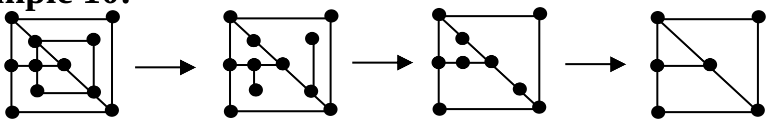

A + B (1) (2) (3) **A**

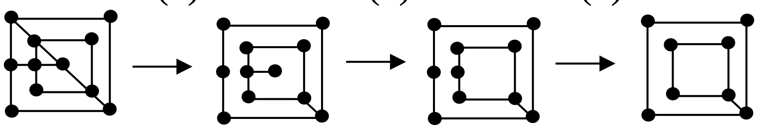

A + B (1) (2) (3) **B**

The **Euler characteristic of a surface** is defined to be the Euler characteristic of any map on that surface.

We can consider an atlas as a topological map, where each map is a face. The edges and vertices of the separate maps are the edges and vertices of the topological map. However in counting them we must take into account the identifications. Each identified pair counts as a single edge. When counting the vertices we must use the identification of the edges to establish which vertices in the atlas represent the same point on the surface. All vertices that are identified together count as just a single vertex. All (connected) surfaces can be represented by a PIE, that is, an atlas with one map. In this case we have F = 1.

Example 11:

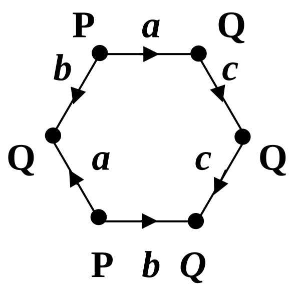

$$\left.\begin{array}{l} V = 2 \\ F = 1 \\ E = 3 \end{array}\right\} \chi = 0$$

Example 12:

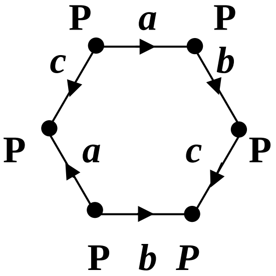

$$\left.\begin{array}{l} V = 1 \\ F = 1 \\ E = 3 \end{array}\right\} \chi = -1$$

Here are the Euler characteristics of the 7 basic surfaces:

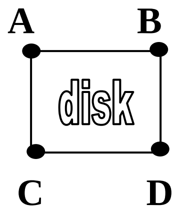 V = 4, F = 1, E = 4 $\chi = 1$	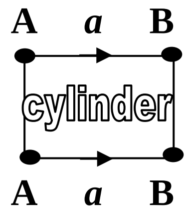 V = 2, F = 1, E = 3 $\chi = 0$	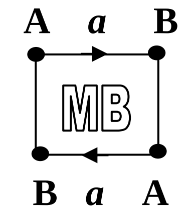 V = 2, F = 1, E = 3 $\chi = 0$	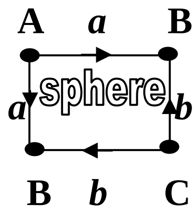 V = 3, F = 1, E = 2 $\chi = 2$

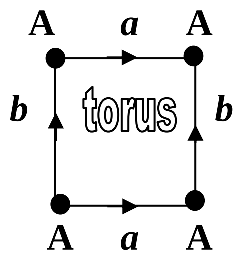 V = 1, F = 1, E = 2 $\chi = 0$	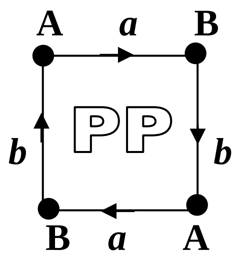 V = 2, F = 1, E = 2 $\chi = 1$	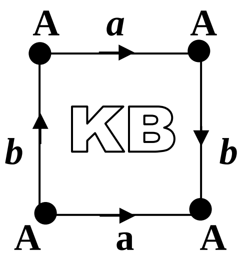 V = 1, F = 1, E = 2 $\chi = 0$

The Euler characteristic is clearly a topological invariant. Surfaces that are homeomorphic to one another have the same Euler characteristic. This is because a homeomorphism from a surface S to a surface T takes any map on S to a map on T with the same number of vertices, faces and edges.

But the converse isn't true. The Euler characteristic doesn't characterise a surface completely – not by itself. A cylinder, a Möbius Band, a torus and a Klein Bottle all have Euler characteristic zero. But they're different surfaces. We need the other invariants.

	CYLINDER	**MB**	**TORUS**	**KB**
orientable?	YES	NO	YES	NO
# holes	2	1	0	0

§4.4. The Weight of a Surface

Let's now tackle the problem of determining the relation between the Euler characteristic of two surfaces and their sum. This will enable us to calculate the Euler characteristic of $m\mathbf{D} + n\mathbf{T}$ or $m\mathbf{D} + n\mathbf{P}$ in terms of m and n.

Such a calculation would be very simple if it was the case that $\chi(\mathrm{S} + \mathrm{T}) = \chi(\mathrm{S}) + \chi(\mathrm{T})$. But clearly this isn't so, since:

$\chi(\mathbf{D}) = 1$ while $\chi(\mathbf{D} + \mathbf{D}) = \chi(\text{cylinder}) = 0$, not 2.

I define the **weight of a surface** S to be

$$\boldsymbol{\chi_0(S)} = 2 - \chi(S)$$

The significance of weight is that it satisfies the simple property that makes calculations easy. The weight of a sum is the sum of the weights.

Theorem 5: $\chi_0(S_1 + S_2) = \chi_0(S_1) + \chi_0(S_2)$.

Proof: Take a map on each of S_1 and S_2 in such a way that one face in each map has three edges. Now cut out these faces and identify the edges of one of these 'triangles' with those of the other. The surface that results is homeomorphic to $S_1 + S_2$.

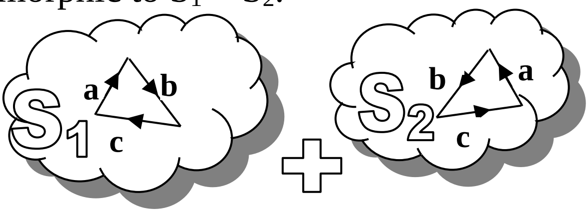

Suppose that the map on S_1 has V_1 vertices, F_1 faces and E_1 edges, and define V_2, F_2 and E_2 similarly. Let V, F and E denote the numbers of vertices, faces and edges of this combined surface $S_1 + S_2$.

Then $V = V_1 + V_2 - 3$ since we've identified the 6 vertices in the two triangular holes into 3 identified pairs. Similarly $E = E_1 + E_2 - 3$.

Finally $F = F_1 + F_2 - 2$ since we lose the two triangular faces.

So $\chi_0(S_1 + S_2) = 2 - \chi(S_1 + S_2)$

$= 2 - (V + F - E)$

$= 2 - (V_1 + V_2 - 3) - (F_1 + F_2 - 2) + (E_1 + E_2 - 3)$

$= 2 - (V_1 + F_1 - E_1) + 2 - (V_2 + F_2 - E_2)$

$= \chi_0(S_1) + \chi_0(S_2)$.

	$\mathbf{S_1}$	$\mathbf{S_2}$	$\mathbf{S_1 + S_2}$
#verts	V_1	V_2	$V_1 + V_2 - 3$
#faces	F_1	F_2	$F_1 + F_2 - 2$
#edges	E_1	E_2	$E_1 + E_2 - 3$
$\boldsymbol{\chi}$	$V_1 + F_1 - E_1$	$V_2 + F_2 - E_2$	$\chi(S_1) + \chi(S_2) - 2$
$\boldsymbol{\chi_0}$	$2 - \chi(S_1)$	$2 - \chi(S_2)$	$\chi_0(S_1) + \chi(S_2)$

Corollary: $\chi_0(m\mathbf{D} + n\mathbf{T}) = m\chi_0(\mathbf{D}) + n\chi_0(\mathbf{T}) = m + 2n$ and
$\chi_0(m\mathbf{D} + n\mathbf{P}) = m\chi_0(\mathbf{D}) + n\chi_0(\mathbf{P}) = m + n$.

The orientability, Euler characteristic and number of holes, completely determine a surface up to homeomorphism.

Theorem 6: If two surfaces, S_1 and S_2, have the same Euler characteristic, same number of holes and are either both orientable or non-orientable then they are homeomorphic.

Proof: If they're orientable, with the same number of holes then $S_1 \approx m\mathbf{D} + n_1\mathbf{T}$ and $S_2 \approx m\mathbf{D} + n_2\mathbf{T}$ for some m, n_1, n_2. If they have the same Euler characteristic, they have the same weight, and so $\chi_0(S_1) = m + 2n_1 = \chi_0(S_2) = m + 2n_2$ whence $n_1 = n_2$.

Similarly, if both surfaces are non-orientable, with the same number of holes, then $S_1 \approx m\mathrm{D} + n_1\mathrm{P}$ and $S_2 \approx m\mathbf{D} + n_2\mathbf{P}$ for some m, n_1, n_2. If they have the same Euler characteristic then we can prove that $n_1 = n_2$ as above.

Example 13: Identify the surface:

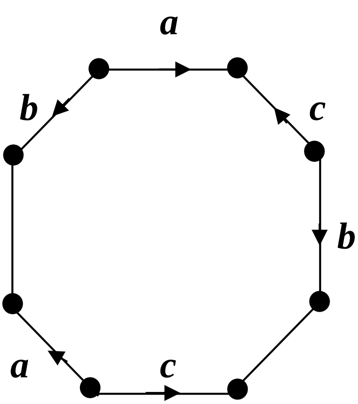

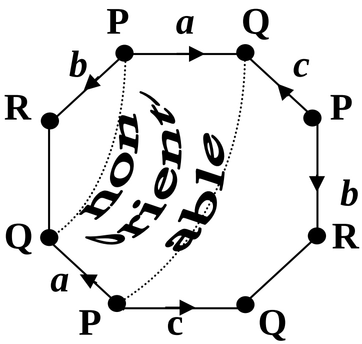 (contains a Möbius Band)	1 hole: R → Q → R V = 3, F = 1, E = 5 $\chi = -1$, $\chi_0 = 3$. SURFACE ≈ **D** + n**P** (for some n) $\chi_0 = 1 + n = 3$ so $n = 2$ SURFACE ≈ **D** + 2**P** = KB with one hole

Example 14: Identify the surface:

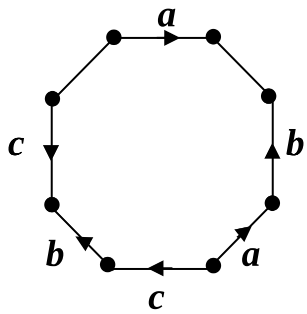

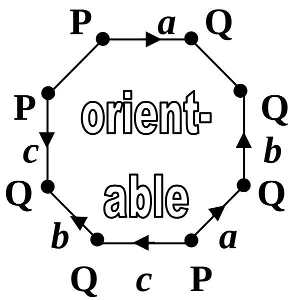	2 holes: P $\to$ P, Q $\to$ Q V = 2 F = 1 E = 5 $\chi = -2$, $\chi_0 = 4$. SURFACE $\approx$ 2**D** + 2**T** = torus with 2 holes

Example 14: Identify the following surface:

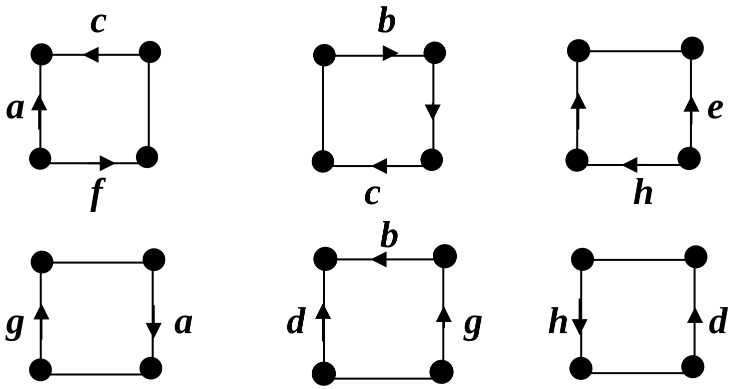

Solution:

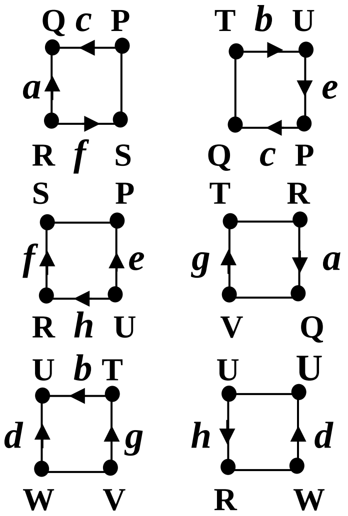	3 holes: P → S → P, U → U and T → Q → V → W → R → T V = 8 F = 6 E = 16 $\chi = -2$ $\chi_0 = 4$.
IF ORIENTABLE: SURFACE ≈ 3**D** + *n***T** (for some *n*) $\chi_0 = 3 + 2n = 4$, no solution.	IF NON-ORIENTABLE: SURFACE ≈ 3**D** + *n***P** (for some *n*) $\chi_0 = 3 + n = 4$, so $n = 1$. SURFACE ≈ 3**D** + **P** ≈ MB + 2 holes

EXERCISES FOR CHAPTER 4

Exercise 1:

A surface, S, is represented by the following polygon, some of whose edges are identified.

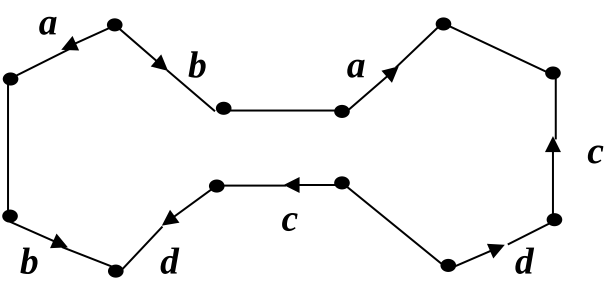

(i) Find the Euler Characteristic of S.
(ii) Is S orientable? Give reasons.
(iii) Express S, up to homeomorphism, in the form

$$m\mathbf{D} + n\mathbf{T} \text{ or } m\mathbf{D} + n\mathbf{P}$$

where **D** represents a disk, **T** a torus and **P** a projective plane.

Exercise 2: The following polygon with identified edges represents a surface S.

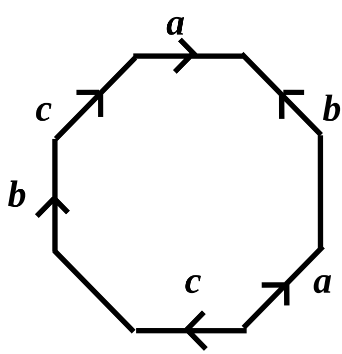

(i) Label the vertices so that vertices have the same label if and only if they represent the same point on S.
(ii) Find the Euler Characteristic of S.
(iii) Is S orientable? Give reasons.
(iv) Express S in the form $m\mathbf{D} + n\mathbf{T} + r\mathbf{P}$.
(v) Are the values of m, n and r unique?
(vi) Carry out a single piece of surgery (one cut and a subsequent joining along one identified pair of edges) to obtain a polygon with identified edges which also represents *S* but which has an *adjacent* like pair.

Exercise 3:
A PIE (polygon with identified edges) can be represented as a word in which each symbol, ***x***, occurs at most twice, either as an *x* and an x^{-1} or as two copies of ***x***. Symbols which occur only once represent unidentified edges and those that occur in a pair represent a pair of identified edges with ***x*** and $\boldsymbol{x}^{-1}$ indicating an unlike pair and a couple of x's indicating a like pair. There is no significance in where the sequence begins. So a torus could be written as $\boldsymbol{aba^{-1}b^{-1}}$ or $\boldsymbol{ba^{-1}b^{-1}a}$ and a Möbius Band as ***abac***, ***baca*** etc.

For each of the following draw the corresponding polygon with identified edges and use surgery (follow the steps in the proof of the Fundamental Theorem on Surfaces) to express it in the form $m\mathbf{T} + n\mathbf{D}$ or $m\mathbf{P} + n\mathbf{D}$ where $\mathbf{D}$ represents a disk, $\mathbf{T}$ represents a torus and $\mathbf{P}$ represents a projective plane.

(a) $\boldsymbol{abca^{-1}dedb^{-1}}$;
(b) $\boldsymbol{abca^{-1}ded^{-1}b^{-1}}$;
(c) ***continuous***.

SOLUTIONS FOR CHAPTER 4

Exercise 1:

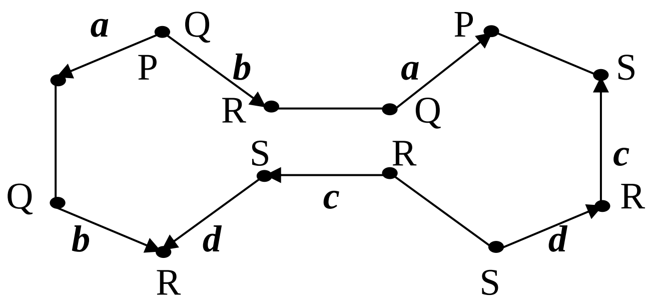

V = 4, F = 1, E = 8 so $\chi = -3$

(ii) S is orientable since there are no like arrows.

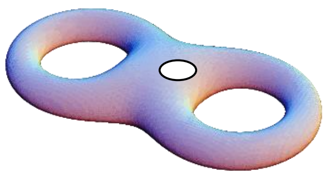

(iii) Since S is orientable
$S \approx m\mathbf{D} + n\mathbf{T}$.
There is one hole: $P \to Q \to R \to S \to P$. So $m = 1$.
$\chi_0 = 2 - (-3) = 5$. Hence $n = (5 - 1)/2 = 2$.
Thus $S \approx \mathbf{D} + 2\mathbf{T}$.
This is a double torus with one hole.

Exercise 2: (i)

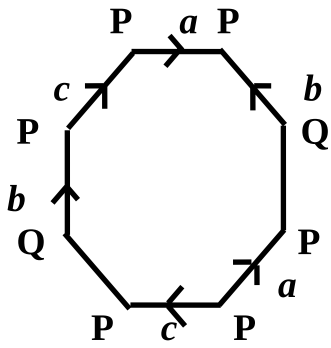

(ii) The Euler characteristic of the surface is

$$\chi = V + F - E = 2 + 1 - 5 = -2$$

(iii) S is not orientable because the polygon with identified edges contains a pair of like arrows (the ***c***'s).

(iv) There is one hole (P→Q→P), and since *S* is non-orientable, S ≈ **D** + *n***P** for some *n*.
Thus its weight is *n* + 1.
Since its Euler characteristic is −2 its weight is
2 – (–2) = 4. Thus *n* + 1 = 4 and so *n* = 3.
Hence *S* ≈ **D** + 3**P**.

(v) Since 3**P** ≈ **T** + **P**, *S* can also be written as
D + **T** + 2**P**. In other words it is the connected sum of a torus with one hole and a Klein Bottle.

(vi)

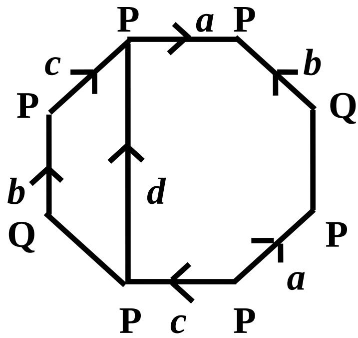

Cut along the edge marked ***d*** and join the pieces by joining along ***c***.

This results in a polygon with identified edges with a pair of adjacent like arrows.

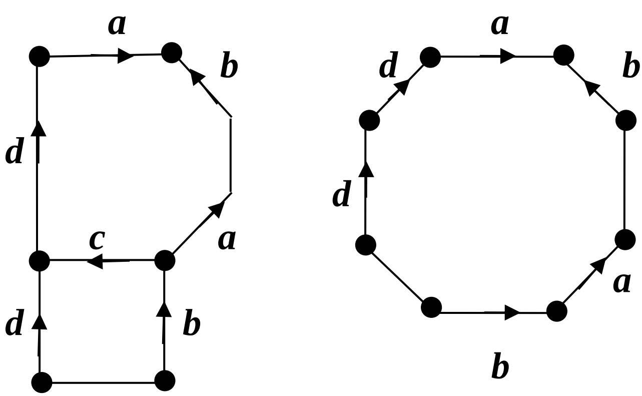

Exercise 3: (To save space the diagrams are omitted).

(a) (iii)V = 2, F = 1, E = 5, $\chi = -2$.

(iv) the surface is non-orientable (***d***'s are a like pair)

(v) #holes = 1

(vi) $\chi_0 = 4$ so the surface is

3**P** + **D**.

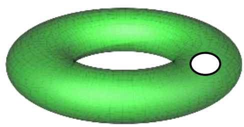

(b) (iii) V = 2, F = 1, E = 5, $\chi = -2$.

(iv) the surface is orientable (no like pairs)
(v) #holes = 2 (they are A→A and B→B)
(vi) χ_0 = 4 so the surface is **T** + 2**D**.

(c) The string ***continuous*** was to be read as a string of letters just as in (a) and (b) with 3 pairs of identified edges (***o***, ***n***, ***u***) and 4 unidentified edges (***c***, ***t***, ***i***, ***s***). Consecutive unidentified edges can be combined giving:

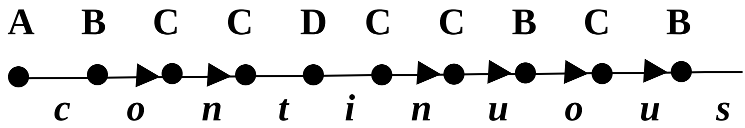

(iii) V = 4, F = 1, E = 7, χ = –2.
(iv) the surface is non-orientable (3 like pairs)
(v) #holes = 2 (they are B→A→B and C→D→C)
(vi) χ_0 = 4 so surface is 2**P** + 2**D**.